Eyeing your Exposure: Quantifying and Controlling Information Sharing for Improved Privacy

Roman Schlegel and Apu Kapadia
School of Informatics and Computing
Indiana University Bloomington
Bloomington, IN, USA
{schleger, kapadia}@indiana.edu

Adam J. Lee
Department of Computer Science
University of Pittsburgh
Pittsburgh, PA, USA
adamlee@cs.pitt.edu

ABSTRACT

A large body of research has focused on *disclosure policies* for controlling information release in social sharing (e.g., location-based) applications. However, less work has considered how *exposed* these policies actually leave users; i.e., *to what extent* are disclosures in compliance with these policies actually being made? For instance, consider a disclosure policy granting Alice's coworkers access to her location during work hours. Alice might feel that this policy appropriately controls her exposure, but may feel differently if she learned that her boss was accessing her location every 5 minutes. In addition to specifying *who* has access to personal information, users need a way to quantify, interpret, and control the extent to which this data is shared.

We propose and evaluate an intuitive mechanism for summarizing and controlling a user's exposure on smartphone-based platforms. Our approach uses the visual metaphor of *eyes* appearing and growing in size on the home screen; the rate at which these eyes grow depends on the number of accesses granted for a user's location, and the type of person (e.g., family vs. friend) making these accesses. This approach gives users an accurate and ambient sense of their exposure and helps them take actions to limit their exposure, all without explicitly identifying the social contacts making requests. Through two systematic user studies ($N = 43, 41$) we show that our interface is indeed effective at summarizing complex exposure information and provides comparable information to a more cumbersome interface presenting more detailed information.

Categories and Subject Descriptors

D.2.2 [**Software Engineering**]: Design Tools and Techniques—*User interfaces*; H.5.2 [**Information Interfaces and Presentation**]: User Interfaces—*Evaluation/methodology*; K.4.1 [**Computers and Society**]: Public Policy Issues—*Privacy*

Symposium on Usable Privacy and Security (SOUPS) 2011, July 20–22, 2011, Pittsburgh, PA USA

General Terms

Design, Human Factors

Keywords

Privacy, location sharing, exposure feedback

1. INTRODUCTION

With the advent and proliferation of mobile devices such as smartphones, more and more people are sharing or broadcasting personal *contextual* information using social-networking services such as Facebook[1] and Twitter.[2] For example, people are now explicitly sharing potentially-sensitive information like their current location—perhaps gleaned from continuous monitoring, explicit "check-ins," (as in Foursquare[3]) or geo-tagged photographs—as well as implicit activity information (e.g., "walking", "running", or "dancing") as deduced from onboard sensors such as accelerometers [23]. In the near future, medical sensors will enable applications such as distributed health monitoring and health-status sharing (e.g., between senior citizens and family) [18]. With such a wealth of personal contextual information always available via mobile devices, it is imperative to protect the privacy of the individual by allowing for control over the dissemination, utilization, and attribution of this data.

At first blush, it seems plausible that the extensive body of research literature within the domain of access control and context sharing could be applied to address such privacy issues. Informally, access control mechanisms allow users to specify *disclosure policies* identifying—either implicitly or explicitly—the sets of principals to whom their data should be made available. Over the years, solutions have been developed to control information release based on the identity of the querier, such as in [9] or in OpenID,[4] the functional roles [28] or other attributes [32] ascribed to the querier within his or her organization, collections of attributes and certifications managed by third parties external to the access control process [34], and dynamically generated proofs of authorization that allow various types of delegation of control [4, 15, 21]. Researchers have also specifically looked at various access-control mechanisms for sharing location [2,10,24] and context [3,12,16,19,33], as well as *spatio-*

[1] `http://www.facebook.com`
[2] `http://www.twitter.com`
[3] `https://foursquare.com`
[4] `http://openid.net`

temporal blurring [7, 8, 11] approaches to reduce the fidelity of shared information for increased privacy.

Although these types of access-control solutions are a vital part of today's computational landscape, the process of crafting disclosure policies is only one aspect of attaining adequate privacy. With regard to context sharing, disclosure policies authorize *queriers* to access a user's information. However, these solutions do not provide the *user* with feedback about the queries; i.e., users do not know to whom and to what extent they are exposed to (authorized) queriers. In our work, we seek to address this larger notion of *exposure* within which pervasive and mobile data sharing must be considered. We aim to provide users with a way to quantify, interpret, and control the extent to which their personal data is accessed, cross-correlated, and disseminated to queriers. We now provide a motivating scenario to illustrate the dynamism of information exchange and the need for suitable exposure control to better manage end-user privacy. In Section 6.2 we provide data from a user study, which shows that users indeed have such concerns about their exposure.

Motivating Scenario: Friend Finders. Location-based services like Google Latitude,[5] Loopt,[6] and Foursquare allow individuals to locate nearby friends. In the context of such a system, one can envision a user, Alice, who sets a disclosure policy enabling her co-workers to access her physical location during work hours to facilitate in-person meetings. This policy may give Alice a baseline perception of exposure control, as it restricts access to her location to a subset of her social connections. However, Alice would likely feel over-exposed if, e.g., she learned that her boss typically monitored her location every five minutes, or that several members of her project team accessed her location while she was visiting a medical specialist during work hours (a situation she had previously not anticipated).

Ideally, these types of applications would provide Alice with some *feedback about the queriers* to help her understand how well her perceived data sharing policy matches reality. For example, her phone might vibrate for increasingly-long periods of time depending on the frequency with which her location is shared. This problem gives way to a rich design space with many challenges. For instance, how can we provide Alice with this sense of exposure awareness without constantly interrupting her? How can Alice react to increased exposure? What level of awareness is appropriate to protect the privacy of the queriers? For instance, Bob might feel uncomfortable if his identity is revealed each time he requests Alice's location; in this case, would simply notifying Alice that "a coworker is accessing your location" be sufficient? As this example motivates, *we need a mechanism to provide Alice with unobtrusive feedback about her exposure while also protecting the privacy of her queriers.*

Our Contributions. As illustrated by the preceding example, the use of existing and emerging social networking applications without an appropriate means of exposure feedback and control can lead to unintended privacy breaches. Since the amount and sensitivity of information being shared using these types of applications is only increasing, it is important that this problem be addressed. In this paper, we take a systematic approach towards studying the exposure problem and developing an exposure control solution. To this end, we make the following contributions:

1. We propose a smartphone-based ambient interface that uses the visual metaphor of *eyes* to provide users with feedback about their exposure. Users can interact with our interface for controlling and limiting their exposure through an intuitive and quick mechanism, as opposed to having to edit cumbersome policies or privacy settings. This interface masks the relationship between *queriers* and *requests*, thereby providing some level of querier privacy.

2. We study peoples' attitudes towards exposure. In our first user study ($N = 43$) we measure how often subjects are willing to share their contextual information with various categories of queriers (i.e., significant other, family, close friends, high-school friends, strangers). We found that subjects differentiated between these categories and for each category, subjects had a concrete number of acceptable accesses per day (e.g., 1 to 2 accesses for high-school friends per day).

3. Informed by our first study, we perform a detailed and systematic evaluation of our visual metaphor via a second user study. For this study, we created a generalized framework for investigating and evaluating various exposure interfaces. Specifically, we developed a point-based game to measure the effectiveness of our interface at conveying exposure information. In this game, the points gained or lost per location-sharing transaction are based on utility functions derived from the data collected during the first study. The framework supports mechanisms by which subjects can limit accesses by queriers, reacting to their exposure. One of our contributions thus is a systematic game-based design template for measuring the effectiveness of exposure interfaces.

4. Results from a second user study ($N = 41$) show that subjects using our "eyes" feedback mechanism are able to obtain equivalent utility when compared to a control group with detailed (but more cumbersome) exposure feedback. Furthermore, results show that our ambient interface is intuitive, and subjects were able to interpret this information more quickly than with the detailed information interface. As such, our ambient interface provides an effective and intuitive means of exposure feedback and control.

Paper outline. We first give an overview of related work in Section 2. We then describe our system model in Section 3 and Section 4 gives an overview of the results of the preliminary user study. In Section 5 we present our smartphone-based ambient interface, followed by the detailed design of our second user study in Section 6. We then present the results of our second study in Section 7, proceed to the discussion in Section 8 and finally conclude in Section 9.

2. RELATED WORK

The problem of providing feedback to users whenever personal or sensitive information is accessed, was studied by

[5] `http://www.google.com/latitude`
[6] `http://www.loopt.com`

Bellotti and Sellen [1] in 1993 in the context of EuroPARC's RAVE system [6]. This system allowed people in the EuroPARC office to share audio or video feeds from their office spaces with colleagues in other offices. Bellotti and Sellen outline a framework that enables a user to limit the dissemination of captured data using basic access controls. Their system also provides very basic feedback when audio/video data is being accessed, e.g., by illuminating an LED next to the camera. While these basic protections are suitable for an environment with a fixed and small user population, they do not provide the level of detail or interaction required for today's mobile environments.

Hsieh et al. [14] introduced *IMBuddy* in a study about control and feedback mechanisms for sensitive information requested via instant messaging (IM) client. They augmented an IM client to provide information about the interruptibility, location, and context (i.e., the current software used) of the user, and made this data available to the user's buddies. Their system allows users to set disclosure policies on a per-buddy basis, and optionally provides feedback to the user about the actual accesses made. This feedback is provided only as a detailed history of which buddies accessed the user's information; this does not protect the privacy of queriers, and requires the user to manually search for and assess excessive exposure conditions. No ambient feedback is provided, and case-by-case or temporary exceptions to policy cannot be made in response to this feedback.

Tsai et al. [31] present a user study of a mobile location sharing application in which one group of users received feedback about accesses made to their location, while another group did not receive any feedback. Their results show that feedback is an important factor for making people feel more comfortable about sharing their location and reducing privacy concerns. In fact, they showed that feedback can actually encourage people to share their location more often. Their study allowed users to allow accesses based on time, although not on a per user basis. In their study, feedback was only provided as a history of accesses, which has to be explicitly searched by users without other immediate or ambient feedback. Their study also did not provide a means of denying accesses on a case-by-case basis.

Patil and Lai [26] carried out a user study to evaluate whether providing feedback while defining initial privacy settings had any effect on how participants chose their initial data sharing settings. Their findings show that giving feedback during the initial definition of sharing settings does not significantly influence participants' data sharing decisions. Their study, however, only focuses on feedback given while defining general privacy settings. Specifically, it does not consider continuous, ambient feedback during the actual use of an information sharing system.

Hong et al. [13] define an architecture for protecting the privacy of users during the disclosure of sensitive information in ubiquitous computing systems. Their design articulates the need for simple control and feedback mechanisms as a basic end-user requirement. Their approach focuses more on access control and information flow control—e.g., allowing emergency services access to precise location of users in case of emergencies—and proposes the use of limited feedback mechanisms like basic notifications. The relationship between ambient feedback mechanisms and disclosure control is unexplored in their approach.

Another interesting approach to immediate feedback regarding disclosed information is presented by Mynatt et al. [25], who develop the idea of *privacy mirrors*. Privacy mirrors not only show what information is collected, but also how frequently and possibly by whom it was accessed. These mirrors are used as separate objects and are not necessarily integrated into context sharing applications or devices, and may violate querier privacy.

A location sharing application by Toch et al. called Locaccino is described in [29,30]. Locaccino allows users to define expressive location sharing rules (e.g., rules that depend on day/date, time, and physical location) for different kinds of queriers. It does not, however, include any continuous or explicit feedback mechanism; users of the service must manually check an audit log and adjust their general preferences if necessary, making it more difficult to get an intuitive grasp of their exposure.

In summary, while previous approaches for mitigating end-user exposure have been proposed, they mainly focus on setting general disclosure preferences and allowing users to revisit them if necessary. Furthermore, most feedback mechanisms described rely on simple histories of accesses; this requires careful examination by the end-user to judge their level of exposure, and may violate querier privacy. In this paper, by contrast, we aim to provide a mechanism for delivering continuous and integrated feedback in an ambient manner. Our approach allows users to control their exposure intuitively without having to revisit general privacy settings explicitly. By masking the relationship between queriers and requests, our approach provides a sense of exposure awareness while helping preserve querier privacy.

3. SYSTEM MODEL

Figure 1 shows a simplified view of the system model that we assume for location (or context) sharing. This system model consists of three main types of entities: *providers*, *data managers*, and *queriers*. We use the term *providers* to refer to the entities in the system who are actively sharing their contextual data with others, and the term *queriers* to denote the principals who request the contextual data of others. A single entity may be both a provider and a querier, and many providers and queriers are likely to exist at any given time. A *data manager* is a logically centralized entity that brokers access to a provider's contextual data. The responsibilities of a data manager include (i) storing the contextual data published by providers, (ii) properly enforcing the disclosure policies maintained by providers, (iii) auditing the actions taken in the system, and (iv) using these audit records to calculate each provider's exposure level.

In the above diagram, solid lines denote administrative tasks, while dashed lines denote interactions resulting from data requests. In particular, the user is assumed to write disclosure policies that are used by the data manager to protect the data pushed by the user. The user will periodically receive exposure updates from the data manager, at which point she may decide to update her policies to better control her exposure. We assume that certain evidence collected during the policy satisfaction process is available only to the data manager, while other portions are available for the provider to use to better understand her current exposure level. For instance, Alice might not learn the identity of the querier who obtained her location, but only that he is one of her friends who is located in the same state.

We note the above model is sufficiently general to model

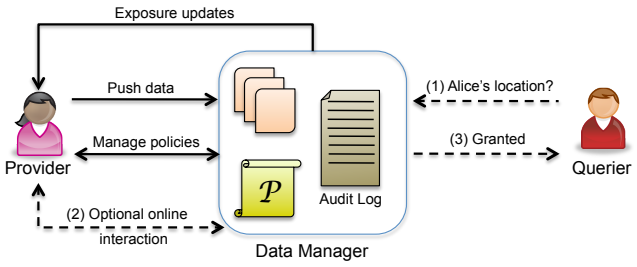

Figure 1: A high-level system model in which information providers leverage a logically centralized data manager to share contextual information with one or more queriers.

a wide range of interesting scenarios. For instance, services like Google Latitude or Facebook Places can be modeled as logically centralized (although physically distributed) data managers that keep track of users' location and status updates, respectively. In this situation, the user must intrinsically trust the data manager to properly enforce her policies and protect her data, which is the dominant assumption in deployed systems today. This model makes it trivial to enforce policies that involve the data manager verifying attributes of the querier that should not be disclosed to the provider. On the other hand, the data provider could also be a physically centralized module built into the mobile device of the provider in situations where data is brokered in a peer-to-peer manner. However, this situation would require the use of cryptographic protocols to make use of private or hidden querier attributes (e.g., [5, 17, 20, 35]). Lastly, we note the data manager need not be trusted to actually view a provider's contextual data, as this data could be stored in an encrypted manner and recovered by the querier using a key that is either pre-issued or revealed during the policy enforcement process.

4. PRELIMINARY STUDY

The location exposure evaluation framework that we discuss in Section 6 requires a baseline notion of tolerable exposure in order to determine the number of points gained or lost due to a particular location access. To calibrate the exposure game described in Section 6, we first ran an initial study to better understand general attitudes towards location exposure. Our study asked a group of 43 participants from Indiana University Bloomington how intrusive they would rate different numbers of requests for their shared personal information by different people, and also asked them how many requests per day they would consider to be acceptable.

For the study, we defined different categories of queriers: *significant others*, *parents*, *close friends*, *high school friends*, and *strangers*. Participants were asked to use a 10-point Likert scale (1 = Very Uncomfortable, 10 = Very Comfortable) to rate their comfort level with various frequencies of location disclosure to a *single* representative from each of the above categories. Figure 2 shows the results of this study (also indicating the standard error of the results), where the X-axis represents the number of accesses made per day by an individual requester, and the Y-axis represents the perceived comfort level of the subject sharing her location.

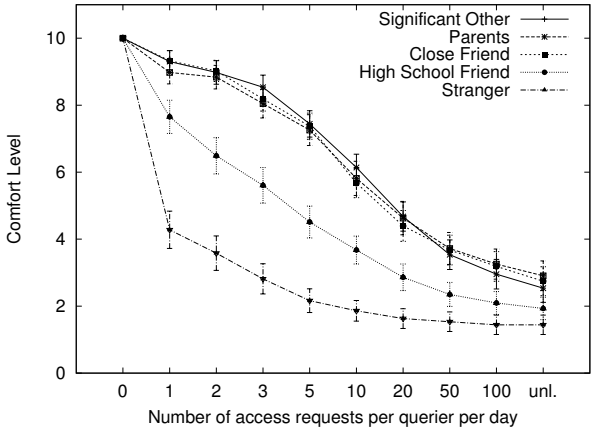

Figure 2: This figure shows how comfortable people felt on a Likert scale of 1 to 10 when sharing their location, depending on the number of accesses made by different people. It illustrates that subjects only distinguished between three categories of queriers and that number of accesses per day have an influence on the comfort level.

Figure 2 suggests that both the type of querier and the frequency of aggregate access matter significantly. One interesting conclusion is that subjects did not distinguish between the categories of significant other, parents, and close friends. Subjects indicated that they are most willing to share information with this collection of requesters, somewhat less willing to share information with high-school friends,[7] and considerably less comfortable when sharing their location with strangers.

The differences in the mean comfort levels are statistically significant (using an independent-samples T test) with $sig < 0.05$ for the range from $x = 1$ request to $x = 50$ requests, except for the difference between high-school friend and strangers at $x = 50$ requests (there is no difference at $x = 0$ as no accesses are made).

We also asked subjects how many times *in total* per hour they would be willing to share their location with any type of queriers. Figure 3 shows their level of comfort for different numbers of aggregate accesses per hour.

Findings. Our initial exploration of people's attitudes towards exposure yielded the following findings to drive the design of an interface to convey and control exposure:

- The type of querier matters and subjects typically distinguished between three different categories of queriers with regard to exposure.
- The number of accesses over time matters. Subjects had diminishing levels of comfort when queriers accessed the location with higher frequencies.
- Static policies like those in Google Latitude and similar services do not accurately capture subjects' attitudes

[7] We note high-school friends are assumed to be alumni, since subjects are adults. If subjects included high-school students, we would expect to see much higher comfort levels for this category.

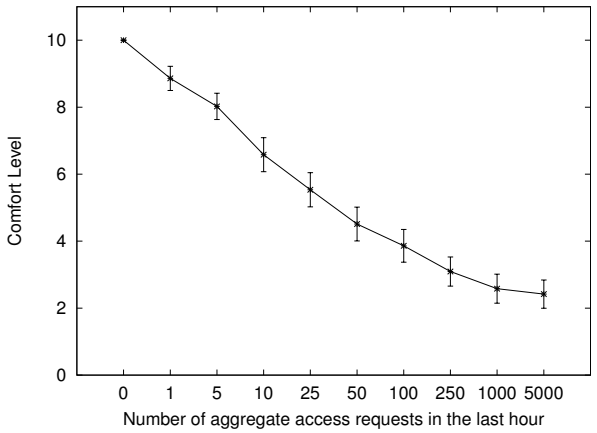

Figure 3: This figure shows how comfortable people felt on a Likert scale of 1 to 10, depending on the number of aggregate accesses made by all queriers in the last hour. As aggregate accesses in the last hour increase, the comfort level of a provider decreases.

towards exposure, and thus an interface to convey and control exposure is needed.

5. AN EYES-BASED INTERFACE

We now describe our smartphone based mechanism for providing unobtrusive and intuitive exposure feedback to users while protecting the privacy of queriers. This design was driven by the findings described in our initial user study, i.e., our application should provide exposure feedback about the number of accesses from different categories queriers, as well as provide a means to control this exposure, which is not found in current static disclosure policies.

5.1 Application overview: The eye metaphor

The central purpose of ambient exposure feedback is to provide intuitive and unobtrusive information to a provider when people access his or her contextual information. Based on the observation that users of location sharing applications typically use their smartphones for numerous activities—e.g., social networking, email, web surfing, and (even) phone calls—we expect users acting as data providers to encounter their home screen several times a day as a natural part of their workflow. As a result, their home screen becomes an ideal location to provide ambient exposure feedback.

Our main idea is to represent accesses by an individual querier as pair of eyes shown in the background of the home screen of a user's smartphone. These eye indicators are thus a metaphor of "being watched" by a person. The presence of a pair of eyes indicates that a querier has accessed the provider's information, while the relative size of those eyes indicates how many times a particular querier has accessed information. A querier making many accesses to the provider's information will thus be represented by a larger pair of eyes than a querier making fewer accesses. As will be described in Section 5.2, the size of these eyes represents a moving average of the number of accesses made and may also be scaled based on the type of querier. For example, accesses made by a high-school friend may result in a faster growth rate than accesses made by a close friend, which would indicate a higher level of exposure.

To help preserve querier privacy, the provider is *not* told the identity of the querier associated with a particular pair of eyes. However, it may be desirable for the provider to infer the category of querier (e.g., "a colleague" vs. "a high-school friend"). In our interface, this may be possible by observing the growth rate of the eyes. The category of querier could also be information provided to the user explicitly by annotations or color variations within the interface (although different colors may raise accessibility issues for people who are color blind); in our study, we deliberately do not provide this extra information in an effort to keep the interface design simple and intuitive. Figure 4 shows a screenshot of our application to illustrate this idea on the Android platform.

Figure 4: An example of ambient exposure feedback. In this example, seven queriers have made accesses and two of them have been blocked temporarily from making further accesses.

5.2 Size of eyes: Signaling exposure

As previously described, a pair of eyes corresponds to some specific querier, whose identity is kept hidden from the provider. The linear size of these eyes (i.e., height and width) primarily depends on the number of accesses made by the querier over the last n hours. The size of a pair of eyes is then given by first setting a certain threshold size, s_{thr}, which indicates how large a pair of eyes should be drawn when the number of accesses reaches an access threshold, a_{thr}, which is deemed critical by the user. When first installing the application, users are asked to answer a questionnaire about how many accesses per day should be permissible for different categories of queriers. For example, the user might pick the access threshold $a_{thr,colleagues} = 5$ for her colleagues, indicating that she is comfortable with any individual colleague accessing her location at most 5 times in a 24-hour timespan. When accesses by a colleague

approaches a_{thr}, the size of the eyes corresponding to that querier approaches s_{thr}.

We also define a minimum eye size, s_{min}, chosen such that a user with even just one access is easily visible on the interface. Given a querier q from the category cat_q who has made some number a_q of requests in the last 24 hours, we calculate the size s_q of the eyes representing q as follows:

$$s_q = s_{min} + (s_{thr} - s_{min}) \cdot \frac{a_q}{a_{thr,cat_q}} \quad (1)$$

Simply put, the size of eyes will vary linearly between the minimum and threshold sizes as the number of accesses varies from 1 to a_{thr,cat_q}, which represents the access threshold for users within the querier's category. The main idea is that regardless of a querier's category, the same size of eyes s_{thr} signals to the provider the same level of exposure *based on their own perceptions of exposure*. If a querier continues making accesses, the size will grow above and beyond the threshold size s_{thr} signaling over-exposure. Users can learn the size s_{thr} during install time, or explicit feedback can be provided when this size is reached. So as to not be overly prescriptive when users should exactly block queriers, we currently rely on users to intuitively react to sizes of these eyes. Our user study shows that users are able to effectively recognize their exposure levels without explicit signals such as color changes or textual feedback.

Also note that the growth rates of a pair of eyes will depend on the threshold a_{thr,cat_q} defined for the category containing this querier. If this value is low (more restrictive), then the eyes will grow to the threshold size relatively fast; if this value is high (more permissive), then the eyes will grow slower. For example, 1 access of a stranger might produce eyes the same size as the eyes for a family member with 10 accesses, to show that the exposure for these two cases is the same, even if the actual number of accesses differ (i.e., revealing information to strangers increases exposure more quickly than revealing information to family members).

While there are several possibilities for summarizing the number of accesses made by a querier in the past n hours, we chose to use a moving average. By using a moving average, we give more weight to recent accesses when measuring exposure. Thus, the contribution of each access to the statistic used during the exposure calculation is adjusted based on when it was made. Specifically, the contribution of a particular access decays by $\frac{1}{n}th$ each hour and the access summary for one particular person is calculated by summing up all of her weighted (decayed) accesses over the last n hours. As an example, with $n = 24$ an access made 3 hours ago would contribute $\frac{21}{24}$ to the total number of accesses in the last 24 hours. In this case the variable a_q in Equation 1 represents the sum of all decayed accesses. Choosing the decay rate as $\frac{1}{24}$ will also ensure that accesses decay naturally to 0 after 24 hours, when they fall out of the sliding window. In future work, we plan to explore other metrics, but this is outside of the scope of our research in this paper.

5.3 Exposure control: Reacting to feedback

Although providing exposure feedback is an important goal, such an interface is only really useful if it allows the user to react to this feedback and better control their exposure. Our interface therefore allows a provider to control his or her exposure by temporarily blocking either individual queriers or all queriers from accessing personal information. Individual queriers who are causing a high level of exposure can be blocked by simply clicking on their eyes and confirming that they should be blocked. Figure 4 shows an example where two queriers have been blocked by the provider for the rest of the day. Given that only the accesses in the last 24 hours contribute to the exposure of a person, a block lasts also 24 hours.

A dedicated button or menu entry can be used to block everybody from accessing information when a user feels that their *aggregate exposure* becomes too high. A provider may determine that even if all the eyes are below the threshold size, that "too many people are watching me" when viewed in aggregate. For the aggregate case, the block is left on until the overall exposure level has dropped to a level the user is comfortable with. At this point, the user may manually remove the block; this is in contrast to individual blocks, which last for 24 hours.

6. STUDY DESIGN

One of our contributions is the design of a game based study to evaluate various exposure feedback and control interfaces (we evaluate two such interfaces in this paper). We now describe this game based design in detail.

6.1 Game-based design

A major goal of this design is to measure the effectiveness at summarizing the accesses made by various queriers without interference from other variables such as individual privacy preferences. Thus we design a game in which subjects are focused on maximizing "points" corresponding to the utility of sharing location and the costs of heightened exposure when location is over shared. In this setup, all subjects play the same game, with their final compensation depending on their score. This type of variable monetary reward incentivizes (rational) users to make an effort to interpret the information summarized by the interface. As a result, we can draw conclusions about the effectiveness of each interface based on the points received.

In the game, subjects are presented with a summary of simulated accesses—as will be described in Section 6.2—over the past hour for a period of 3 simulated days. After each simulated hour, subjects act upon the exposure information provided by the interface (e.g., by selectively blocking queriers that have exceeded their query thresholds) to maximize their score at the end of the study. Sharing one's location information with a querier would always make the user gain 3 points, to reflect that sharing one's location has a utility value. At the same time, to model the situation where sharing one's location information too many times leads to overexposure, there is a certain probability that the provider will also lose 6 points each time they share location information with a querier. The probability of losing points increases with the number of times location information has been shared with that particular querier in the last n (e.g., 24) hours. This function is modeled such that when the number of accesses approaches the threshold number of accesses corresponding to an average comfort level of 5, the expected point gain would be 0. Thus subjects do not have an incentive to keep allowing accesses after this threshold. In our user study, this probability function is modeled on the exposure preferences derived from our initial user study (Section 4) and is described in more detail in Section 6.2. In addition to the points lost per individual querier, the subject

can lose 2 points if the aggregate number of accesses becomes too high based on an aggregate probability loss function.

As mentioned earlier, a points based design focuses the analysis on the effectiveness of the interface to summarize accesses to the subject. Now that we have described the general framework of our user study, in the next subsection we describe the details on how accesses are simulated and how the probabilistic point functions are constructed.

6.2 Details of user study design

To evaluate our two feedback mechanisms, we planned a user study where one group of users would be given the eyes based interface, while a second group (the control group) would be given a more detailed information interface. We first provide details on the detailed information interface, and then describe how queries were simulated, and provide details on the probabilistic point loss functions.

Control: Detailed information interface. In this paper, we specifically evaluate our eyes based interface against a "detailed information" interface, with the control group getting the detailed information interface. This prototype provides feedback by giving exact and detailed information about the accesses made by other people. We designed and created this interface specifically to provide unfiltered, raw feedback to users, because we could not find an existing, comparable feedback interface that provided this kind of information. We expect subjects to perform better (but take longer to digest the information) with the detailed information interface, but we also seek to measure if there is a significant penalty to using the eyes-based interface. We also seek to measure if subjects are able to respond to our interface more quickly compared to the detailed information interface.

The detailed information interface (shown in Figure 5) will, at any given point in time, tell the user which person has accessed information how many times in the last 24 hours. The user can then individually block people or summarily deny access to everybody. To reduce clutter and improve usability, the prototype displays the top 3 people per category with the most accesses in the last 24 hours since those are the queriers the provider is likely to block. The names used such as "Family 9" serve as pseudonyms for "Some family member."

Simulation of queries. As part of the game, subjects have to respond to feedback about accesses. The study simulates hourly accesses over a period of three days, each day from 7am to 8pm (13 hours simulation time per day[8]), during which different queriers make different numbers of accesses to the subject's information. After each simulated hour, the subject is presented with exposure feedback and has the option to take action. The categories of queriers used in the study are *friends/family*, *high-school friends* and *strangers*.

To determine how many accesses should be made in total over the three day period we decided that there would be 22 queriers in total, 12 family/friends, 6 high-school friends and 4 strangers. These numbers are somewhat arbitrary (although, intuitively there should be more family/friends than high-school friends and the smallest numbers of strangers), but the exact numbers are not important for the study, as the point-based game is designed such that it does not depend on the exact composition of queriers. 22 is also a number of queriers which could be comfortably displayed on the screen space available for the feedback interface.

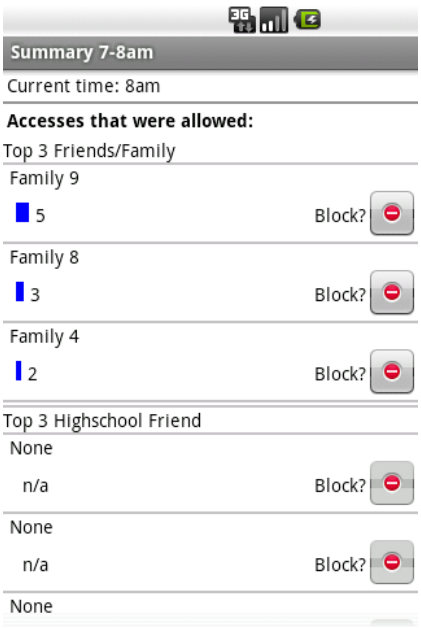

Figure 5: A screenshot of the detailed information feedback interface. For each category of queriers the interface displays the top 3 queriers in terms of accesses and provides options to control exposure by blocking individual queriers (names used such as "Family 9" serve as pseudonyms for "Some family member").

Next, we referred to the results from the preliminary study described in Section 4 to determine the number of accesses each day which would lead to a moderate level of comfort (5 in a scale from 1 to 10). Thus instead of each subject imposing their own privacy preferences, the game is set up to make subjects behave like the average person, who would want to block queriers after the same number of accesses (depending on category). The averages worked out to 13 for family/friends, 2 for high-school friends and 1 for strangers (more on this below). The total number of accesses for each category was then calculated as the number of queriers in a category times the 3 days times the number of accesses corresponding to a moderate level of exposure. This came to 468 for family/friends, 36 for high-school friends and 12 for strangers, over the course of the whole simulation.

Intuitively, if those total accesses were uniformly distributed across all queriers, the provider would be moderately exposed towards each querier, thus requiring some form of exposure control in the study. To make sure that some queriers access the provider's information more often, leading to elevated exposure and thus forcing the provider to take action, we did not distribute the total number of accesses uniformly across queriers. Instead, each category of queriers was divided into three groups: one group with queriers who make many accesses a day, one group with

[8]Because only 13 hours were actually simulated in the study, the decay rate for the calculation of the exposure level was set to $\frac{1}{13}$ for Equation 1, instead of $\frac{1}{24}$.

queriers who make a moderate number of accesses a day, and one group with queriers who make only a few accesses a day. Table 1 shows the distribution of the number of queriers for each group and category combination.

	High	Medium	Low	**Total**
Friends/Family	2	3	7	**12**
High-school Friends	1	2	3	**6**
Strangers	1	1	2	**4**

Table 1: Number of people with different frequency of accesses for different categories of queriers.

Among the different groups, queriers in the group with a high frequency of accesses were 5 times more likely to access information than the group with low frequency, while queriers in the group with medium frequency were 3 times more likely to access information than the group with low frequency. Furthermore, accesses were distributed uniformly over the simulated time span. The numbers 5 times and 3 times are somewhat arbitrary; the important point is that there is a guaranteed variation among accesses, to make sure that the provider is forced to take action to prevent elevated exposure. These parameters taken together ensure that subjects have to react to exposure incidents for different categories of queriers and different groups of queriers who may result in heightened, moderate or low exposure. Again, we stress that the game-based design allows us to make such arbitrary choices to evaluate the relative effectiveness of the two interfaces.

Point loss function. We model the point loss function on the average user as determined by the first user study. This choice ensures that the relative values of the utility functions corresponding to the three categories of queriers are meaningful and the rate of change in eye sizes for example would reflect what an average subject would encounter if the interface were used in real life. Thus, from the preliminary study described in Section 4, we extracted probability functions for losing points in the exposure game. We first fitted curves through the graphs obtained from the preliminary study. Because the results indicated that significant other, parents and close friends are treated very similarly, we combined them into one curve. It turned out that an exponential equation of the following form:

$$y = a \cdot e^{b \cdot x} + c \cdot e^{d \cdot x} + \mathrm{f} \tag{2}$$

Using a, b, c, d and f as the parameters of the curve proved a good compromise between fitting the curves accurately and limiting complexity. All fitted curves are shown in Appendix A.

From the curves indicating how comfortable subjects were with a certain number of accesses, we then extracted a probability function. Specifically, we set a comfortable level of 5 to correspond to a probability of 0.5 of losing points when sharing information. This means that sharing one's location more often than 10 to 15 times per day with a close friend, for example, will result in a probability of more than 0.5 of losing points. Below 10 to 15 times per day the probability of losing points would be lower than 0.5. In mathematical terms, if y is considered the level of comfort, then the probability of losing points is calculated as $1 - \frac{1}{10} \cdot y$. The resulting curves are shown in Appendix A.

Using this probability-based approach for the exposure game, subjects were asked to try and maximize their points, and were incentivized to do so by making their compensation dependent on the number of points obtained. Subjects were guaranteed $5, but could earn up to $8 depending on their final score. Because each access would make a subject gain points, they had an incentive to share their location information. On the other hand, the probability function of losing points cautions them against sharing their location information all the time. These two mechanisms together incentivized the subjects to actually pay attention to the feedback and block individual or all users if they wanted to maximize their points.

6.3 Study Implementation

Both the eyes based interface and the detailed information interface (for the control group) were implemented as separate applications on the Android smartphone platform. Both user study applications (with the detailed information interface corresponding to the control group, and the eyes based interface corresponding to the experimental group) had the following phases:

Introduction 1: Subjects were first shown an introductory screen, explaining that we were going to ask them in a first part to set a general privacy policy they would feel comfortable with. It was also explained that this part is independent of the simulation following in the second part, and that the settings collected are for reference purposes only, they do not influence or play any part in the simulation.

Privacy Settings: Subjects were then asked to decide on the aforementioned privacy settings by filling out a matrix on the screen (see Figure 6).

Introduction 2: The second set of introductory screens focused on explaining the simulation part of the study:

- for the eyes based interface study, subjects were shown the size of eyes corresponding to the midway point where it becomes more likely than not that they would lose points if more accesses were allowed
- for the detailed information interface study, subjects were given specific ranges of number of accesses where it became more likely that they would lose points for further accesses (e.g., 10-15 for family/friends)

Simulation: After the second set of introduction screens, subjects were taken to the simulation part of the study. During the simulation, the application simulates three days (7am to 8pm) of accesses by different queriers. Each hour, the subject is given feedback, either eyes based or detailed information, and then has the chance to act on it, for example by blocking individual queriers. Once the subject is satisfied with his or her response, he or she can click on a button to advance the simulation by one hour.

End: Once 3 days have been simulated, the study ends and displays a screen with the total number of points achieved.

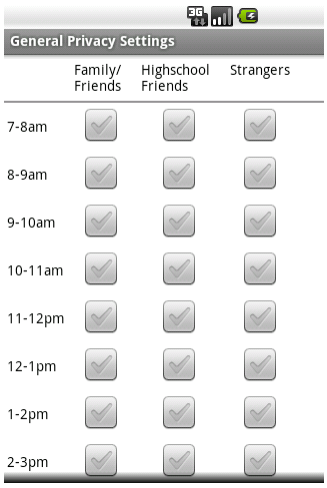

Figure 6: Screenshot of the interface where subjects were asked to define a general privacy policy. The policy allows to define hour-by-hour access rules for three different categories of queriers.

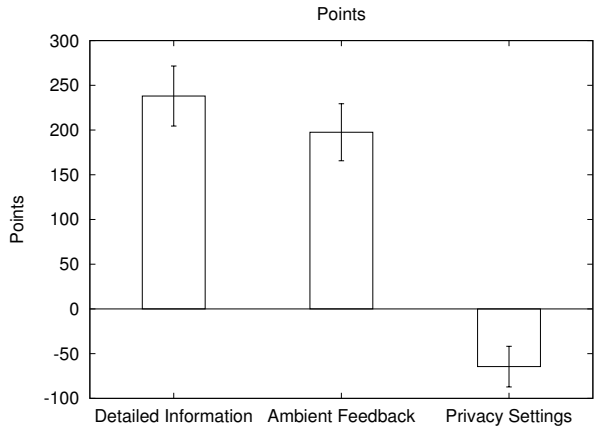

Figure 7: This figure shows the average number of points achieved by subjects. They performed best when given detailed information, but only slightly worse when given ambient feedback. When only using a static privacy policy to allow/deny accesses, the achieved number of points was significantly lower.

6.4 Study Procedures

Recruitment. Subjects were recruited by posting fliers on the Indiana University Bloomington campus, as well as sending out an email to the School of Informatics and Computing undergraduate mailing list. Participants were thus mostly, but not exclusively, undergraduates. The study was open to anybody of age 18 and older with no other conditions.

Subjects were asked to send an email to an address specified on the flier and were then scheduled to come to the Informatics building at IUB and take the study in a room specifically set aside for the study.

Renumeration. Subjects were paid a base amount of $5 for taking part in the study, plus up to $3 in addition to the base amount depending on their performance in the study (i.e., the number of points achieved), for a maximum of $8.

Study. Once subjects arrived at the room set aside for the study at the appointed time, they were given a study information sheet as required by the IRB, detailing the purpose of the study and explaining in more detail their involvement in the study. If a subject had no questions regarding the study, they were then taken to a computer which was running the exposure application within an Android emulator to simulate the look and feel of an Android smartphone.

Because there were two different groups, one with eyes based feedback and one with detailed information feedback, subjects 1,3,5, etc. were given one study and participants 2,4,6, etc. the other. This ensured that we had approximately the same number of people for each interface no matter the final number of subjects.

Once a subject started the study on the assigned machine, he or she was first shown a number of introductory screens giving more details about the exact procedures of the study and explaining how points were gained and lost during the study. Subjects were then left to themselves and unobserved (unless they had any questions regarding the study) until they had reached the end of the study (which was indicated to them on the screen). They were then paid $5 to $8 according to the number of points achieved.

6.5 Ethical considerations

The user studies reported in this paper were approved by the Indiana University IRB.

7. RESULTS

We concluded our user study with 41 participants, in which 21 participants were given detailed information feedback and 20 participants were given eyes based feedback. Apart from recording the number of points achieved by each participant, we also measured the time they took to complete the study (including reading the introductory screens, defining a privacy policy and running the simulation).

7.1 Evaluating effectiveness

On average, subjects taking the detailed information study scored 238 points with a standard error of 33.58. Subjects taking the eyes based study scored on average 197.55 points with a standard error of 31.87 (see Figure 7). As we had expected, the average points scored was lower with the eyes based interface (because it does not provide detailed information about accesses). However, the difference in means between detailed information and ambient feedback is not statistically significant (independent-samples T test, $df = 39, t = 0.87, sig = 0.39$). It is possible that with more subjects this statistical significance would be achieved. Nevertheless, the means are close enough that we conclude the eyes based interface is as effective or almost as effective as the detailed information interface at conveying exposure information to users. To compare these relative means with another baseline, next we examine how many points a subject would have earned solely through static privacy settings (as would be used in existing applications today).

Baseline score using static privacy settings. In addition to the number of points achieved by users, we calculated what the achieved number of points would have been had accesses been permitted or denied based upon the initial disclosure policy specified by the user in the "Privacy Settings" phase of our study. Since the loss of points is determined by a probabilistic process, we run this calculation 1000 times for each user to get an average number of points they would have achieved if accesses were governed by the privacy settings they had set. The result is that, on average, participants would have scored -64.5 points (shown in Figure 7 under "Privacy Settings") with a standard error of 22.69.

Without any exposure feedback, users of social sharing applications have no choice but to rely on such privacy settings to control the exposure of their location As a result, this score reflects how such a static policy would have performed in our game absent any exposure feedback. Given that our point loss functions were developed using the results of the study discussed in Section 4, the disparity in points earned when using either feedback application vs. using a static disclosure policy quantifies the degree to which static disclosure policies fail to address the exposure needs of users. While this score would have varied depending on the simulation parameters, it nevertheless provides a baseline illustrating that subjects actively using our interfaces can significantly improve their exposure level.

7.2 Evaluating ease of use

When measuring the time that participants took to complete the study, we found that there is a significant difference between the two studies. To account for the non-normal distribution of the time measurements, we use a log-transform for the statistics. We found that users required on average 989 seconds ($\log(time) = 2.989013$)[9] for the detailed information study. For the ambient feedback study, on the other hand, participants required only an average of 645 seconds ($\log(time) = 2.790802$) (see Figure 8, which shows the log-transformed values). Furthermore, the difference in mean between the two forms of feedback (detailed information and eyes based feedback) is statistically significant for the log-transform at a 95% confidence level (independent-samples T test, $df = 39, t = 6.329, sig = 0.00000018$). Since the structure of both studies was exactly the same, this difference in time possibly corresponds to the ease with which subjects were able to interpret their exposure information and act upon that information. We thus have some evidence suggesting that the eyes based interface is more intuitive and less cumbersome than the detailed information interface.

7.3 Limitations of the study

Demographic and number of subjects. Because we advertised and recruited subjects on-campus, most of them were college undergraduates and many of them familiar with using smartphones. Ideally a more elaborate study would attempt to measure the effectiveness of the eyes based interface for different demographics. In addition, for both our studies we were able to recruit a total number of approximately 40 subjects each ($N = 43, 41$). A larger number of subjects would provide more statistically significant results.

[9]Technically, when taking the logarithm of a time value the value is first divided by seconds to get a dimensionless value, i.e., $\log(time/s)$, but for easier readability we simply use $\log(time)$.

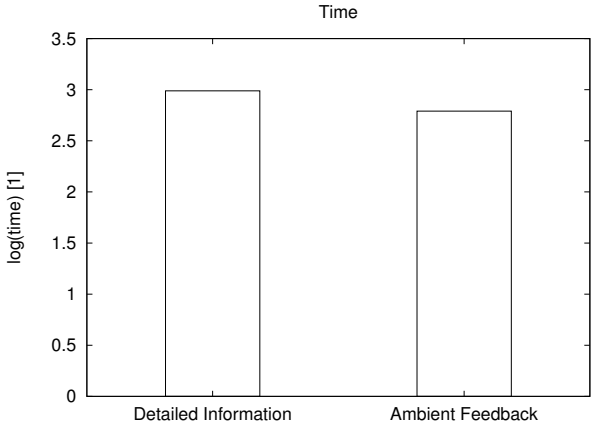

Figure 8: This figure shows the average time taken by participants to complete the study in log(time). Subjects given the detailed information study took significantly longer to complete the study than subjects given the eyes bases study.

Timing. In our results, the time measured for completing the study included the entire study: i.e., the introductory screens, defining a privacy policy and the simulation. In retrospect, it would have been useful to record more precise timing measurements. For instance, recording the time required for completing only the simulation portion of the study, or maybe even the time required for each iteration in the simulation (i.e., for handling the feedback after each simulated hour).

Screen Size. As was noted in Section 6.2, our exposure awareness application begins to become cluttered if a user's location is accessed by many individuals. Although outside of the scope of the current study, it would be interesting to (i) examine variations in the number of queriers with whom providers typically share location information over the course of a given time period, (ii) understand the cognitive workloads associated with managing exposure awareness for various numbers of queriers, and (iii) develop ambient exposure awareness applications that strike a balance between accuracy of information and cognitive workload as the amount/frequency of data sharing increases.

Simulated time. An alternative to a simulated query based study such as ours would be a more elaborate study involving subjects using our application in real location-sharing settings. Nevertheless, we emphasize that this study was designed to evaluate how effective the eyes based interface was at conveying exposure information. Our game-based study was specifically designed to avoid such an elaborate study so that there was no interference from other issues such as individual privacy preferences.

8. DISCUSSION

In this paper we focused on developing an ambient interface for location sharing, while taking into consideration the privacy of queriers. As future work, there are several issues and avenues we are exploring, which we discuss below.

Generalized social context sharing. Projects such as CenceMe [23] provide a platform for users to share various types of personal context with their friends though social networking applications like Facebook. The CenceMe application makes use of the iPhone's onboard microphone, camera, and accelerometer to detect personal context. For example, ambient sound can be used to differentiate work and social settings. The accelerometer can be used to detect various forms of physical activity such as walking, running, or bicycling. We expect such applications to be merged with more mainstream services such as Google Latitude and Loopt, which currently focus on location only. The questions raised with Friend Finders are heightened with context sharing, since location is only *one* type of context. Applications like CenceMe will require that users be able to quantify exposure along several dimensions (corresponding to each type of context) in a usable way. We are exploring ways in which exposure information for various context can be summarized intuitively to users.

Querier privacy. As discussed in Section 2, the findings of Tsai et al. [31] indicate that providing users with feedback about the identities of the individuals querying their location can increase users' levels of comfort and participation in location sharing systems. However, the revelation of query history can also cause unease—and possibly termination of use—for users of social information sharing systems [22]. To balance these competing goals, the interface evaluated in this paper provided queriers with a level of pseudonymous privacy not unlike that provided by techniques like k-anonymity [27]. In particular, our interface allows the user to ascertain that they have been queried by a unique individual, and perhaps infer the user category (e.g., family, friend, etc.) of the querier. There are several avenues for future research on blending querier privacy with exposure feedback. For example, queriers may be allowed a certain number of anonymous accesses before registering on our eyes interface. Furthermore, queriers may be required to reveal their identities beyond a certain number of accesses, and our ambient approach is used between these thresholds. Queriers may have their own privacy preferences and these would need to be evaluated in conjunction with the provider's policies. In our system model this evaluation can be done by the data manager. These scenarios motivate the need for "exposure aware policies" that we discuss next.

Exposure aware policies. In this study subjects were asked to define a privacy policy which allowed them to define hour-by-hour privacy settings for three different categories of people. Although this type of disclosure policy offers more fine-grained controls than many existing approaches, it is in fact quite rudimentary. In our future research, we plan to develop exposure control extensions to established policy architectures that will help users control their exposure and refine policies over time. Interesting features to consider include enabling fine-grained control at the user level, including explicit query rate limits for users or categories of users, and including sharing preferences that are based on physical locations or attributes ascribed to these locations. Furthermore, it would be interesting to consider policies that are "polymorphic" and stateful, in that they change with respect to the user's current exposure, rather than being purely static credential- or identity-based policies. Developing a tight feedback loop that guides the revision of disclosure policies based on a user's reactions to exposure awareness interfaces like the one described in this paper would also be a fruitful area of future work.

Different ways of combining information. In the eyes based interface, we are currently using the size of the eyes as the primary means of conveying exposure information. However, there are many other attributes of this interface that could also be utilized for this purpose. One option would be to use colors to express information about the number of accesses or the type of querier. For instance, a low number of accesses could result in green eyes, a moderate number of accesses would be reflected by orange eyes, and an excessive number of accesses could lead to red eyes. Alternatively, the color of the eyes could be varied to identify them as belonging to a particular class of users (e.g., family, friends, or strangers). Care must be taken, however, to ensure that such modifications would not render the interface useless to color-blind individuals.

More complex social rules. The experiments run in this study represent a first step in characterizing user perceptions of exposure. We focused on a simple model of interaction based on the user/querier relationship and the query frequency. In the real world, however, social interactions are usually more complex. For example, a user might not be willing to share his or her contextual information with all family members equally; e.g., Alice might be comfortable sharing her location with one of her siblings, but not with her parents. Likewise, a user's current location is also likely to influence her willingness to share, perhaps in combination with the identity or role of the requester. For instance, Alice might be willing to share "office" locations with co-workers, but not want to share her "out of office" activities with these same co-workers. Thus exposure feedback would need to incorporate additional information to provide users with a better sense of their privacy. These more complex dimensions will need to be explored to fully realize the potential of exposure control interfaces.

Different modes of ambient feedback. The ambient feedback through eyes shown on the home screen is an unobtrusive mechanism which relies on a user periodically checking the home screen, maybe every couple of hours. There are other mechanisms for ambient feedback, for example, feedback through vibrating the phone when accesses are made, or periodically as a function of aggregate accesses. Some smartphones also have notification LEDs, which could be used to give feedback on the current exposure level, for example by letting the LED blink more quickly as the exposure level increases. Future work could address such modes of ambient feedback on smartphones.

9. CONCLUSIONS

With the advent of social networking applications, users are freely sharing more information than ever before, and doing so without the benefit of feedback telling them how often and to what extent this information is being accessed. This mismatch between a user's conception of how often his or her information is accessed and the actual pattern of accesses can lead to a loss of privacy. While a large body of work has focused on disclosure policies, little attention has been given to the problem of user exposure: i.e., the extent to which disclosures *in compliance with* an access control policy are made. Findings from the first of two studies discussed in this paper show that user perceptions of exposure depend both on inter-personal relationships between the querier and provider, as well as the frequency of access. This implies that existing approaches that depend largely on static disclosure policies are inappropriate for mitigating exposure threats. We then present an unobtrusive and intuitive ambient interface based on the metaphor of eyes for providing exposure feedback and reactive sharing controls to users of context sharing applications. Through a detailed user study we show that this interface is effective at conveying meaningful exposure information to users, and is easier to use than approaches that depend on detailed access histories. This type of feedback on access patterns enables users to better reign in excessive sharing and improve their privacy. We believe further research on exposure feedback and control, as outlined in Section 8, is needed to fully realize the potential of context sharing applications.

10. ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Awards CNS-1016603 and CNS-1017229.

We thank Vaibhav Garg, Sameer Patil, Yann Le Gall, Naveed Alam, and the anonymous reviewers for their helpful comments. We also thank Xiaoyong Zhou and Rui Wang for testing our user study.

11. REFERENCES

[1] V. Bellotti and A. Sellen. Design for privacy in ubiquitous computing environments. In *ECSCW*, pages 75–, 1993.

[2] J. R. Cuellar, J. B. Morris Jr, D. K. Mulligan, J. Peterson, and J. M. Polk. Geopriv requirements, Feb. 2004. RFC 3693.

[3] A. K. Dey. *Providing Architectural Support for Building Context-Aware Applications*. PhD thesis, College of Computing, Georgia Institute of Technology, Dec. 2000.

[4] C. M. Ellison, B. Frantz, B. Lampson, R. Rivest, B. Thomas, and T. Ylonen. SPKI certificate theory. IETF Request for Comments RFC 2693, Sept. 1999.

[5] K. Frikken, M. Atallah, and J. Li. Attribute-based access control with hidden policies and hidden credentials. *IEEE Transactions on Computers*, 55(10):1259–1270, 2006.

[6] W. W. Gaver, T. P. Moran, A. MacLean, L. Lövstrand, P. Dourish, K. Carter, and W. Buxton. Realizing a video environment: Europarc's rave system. In *CHI*, pages 27–35, 1992.

[7] B. Gedik and L. Liu. Location privacy in mobile systems: A personalized anonymization model. In *Proceedings of the 25th IEEE International Conference on Distributed Computing Systems*, pages 620–629, Columbus, OH, USA, June 2005.

[8] M. Gruteser and D. Grunwald. Anonymous usage of location-based services through spatial and temporal cloaking. In *The International Conference on Mobile Systems, Applications, and Services (MobiSys)*, pages 31–42, May 2003.

[9] M. A. Harrison, W. L. Ruzzo, and J. D. Ullman. Protection in operating systems. *Commun. ACM*, 19(8):461–471, 1976.

[10] U. Hengartner and P. Steenkiste. Protecting access to people location information. In *Proceedings of the First International Conference on Security in Pervasive Computing*, pages 25–38, Boppard, Germany, Mar. 2003.

[11] B. Hoh and M. Gruteser. Protecting location privacy through path confusion. In *SECURECOMM '05: Proceedings of the First International Conference on Security and Privacy for Emerging Areas in Communications Networks*, pages 194–205. IEEE Computer Society, 2005.

[12] J. I. Hong and J. A. Landay. An architecture for privacy-sensitive ubiquitous computing. In *Proceedings of MobiSys 2004*, pages 177–189, Boston, MA, USA, June 2004.

[13] J. I. Hong and J. A. Landay. An architecture for privacy-sensitive ubiquitous computing. In *MobiSys*. USENIX, 2004.

[14] G. Hsieh, K. P. Tang, W. Y. Low, and J. I. Hong. Field deployment of *imbuddy* : A study of privacy control and feedback mechanisms for contextual im. In J. Krumm, G. D. Abowd, A. Seneviratne, and T. Strang, editors, *Ubicomp*, volume 4717 of *Lecture Notes in Computer Science*, pages 91–108. Springer, 2007.

[15] T. Jim. SD3: A trust management system with certified evaluation. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 106–115, May 2001.

[16] A. Kapadia, T. Henderson, J. J. Fielding, and D. Kotz. Virtual Walls: Protecting Digital Privacy in Pervasive Environments. In *Proceedings of the Fifth International Conference on Pervasive Computing (Pervasive)*, volume 4480 of *LNCS*, pages 162–179. Springer-Verlag, May 2007.

[17] A. Kapadia, P. P. Tsang, and S. W. Smith. Attribute-Based Publishing with Hidden Credentials and Hidden Policies. In *Proceedings of The 14th Annual Network and Distributed System Security Symposium (NDSS)*, pages 179–192, Mar. 2007.

[18] D. Kotz, S. Avancha, and A. Baxi. A privacy framework for mobile health and home-care systems. In *Workshop on Security and Privacy in Medical and Home-Care Systems (SPIMACS)*. ACM Press, November 2009.

[19] M. Langheinrich. A privacy awareness system for ubiquitous computing environments. In *Proceedings of UbiComp 2002*, pages 237–245, Göteborg, Sweden, Sept. 2002.

[20] J. Li and N. Li. OACerts: Oblivious attribute based certificates. *IEEE Transactions on Dependable and Secure Computing*, 3(4):340–352, Oct. 2006.

[21] N. Li and J. C. Mitchell. RT: A role-based trust-management framework. In *Proceedings of the Third DARPA Information Survivability Conference and Exposition (DISCEX III)*, pages 201–212, Apr. 2003.

[22] R. Metz. Friendster outs voyeurs, Oct 2005. `http://www.wired.com/culture/lifestyle/news/2005/10/69106`.

[23] E. Miluzzo, N. D. Lane, K. Fodor, R. Peterson, H. Lu, M. Musolesi, S. B. Eisenman, X. Zheng, and A. T. Campbell. Sensing meets mobile social networks: the design, implementation and evaluation of the CenceMe application. In *SenSys '08: Proceedings of the 6th ACM conference on Embedded network sensor systems*, pages 337–350, New York, NY, USA, 2008. ACM.

[24] G. Myles, A. Friday, and N. Davies. Preserving privacy in environments with location-based applications. *IEEE Pervasive Computing*, 2(1):56–64, Jan.-Mar. 2003.

[25] E. Mynatt and D. Nguyen. Making ubiquitous computing visible. In *Proceedings of the 2001 CHI Conference on Human Factors in Computing Systems*, 2001.

[26] S. Patil and J. Lai. Who gets to know what when: configuring privacy permissions in an awareness application. In G. C. van der Veer and C. Gale, editors, *CHI*, pages 101–110. ACM, 2005.

[27] P. Samarati. Protecting respondents' identities in microdata release. *IEEE Transactions on Knowledge and Data Engineering*, 13(6):1010–1027, 2001.

[28] R. Sandhu, E. J. Coyne, H. L. Feinstein, and C. E. Youman. Role-based access control models. *IEEE Computer*, 29(2):38–47, Feb. 1996.

[29] E. Toch, J. Cranshaw, P. H. Drielsma, J. Y. Tsai, P. G. Kelley, J. Springfield, L. F. Cranor, J. I. Hong, and N. M. Sadeh. Empirical models of privacy in location sharing. In J. E. Bardram, M. Langheinrich, K. N. Truong, and P. Nixon, editors, *UbiComp*, ACM International Conference Proceeding Series, pages 129–138. ACM, 2010.

[30] E. Toch, J. Cranshaw, P. Hankes-Drielsma, J. Springfield, P. G. Kelley, L. Cranor, J. Hong, and N. Sadeh. Locaccino: a privacy-centric location sharing application. In *Proceedings of the 12th ACM international conference adjunct papers on Ubiquitous computing*, Ubicomp '10, pages 381–382, New York, NY, USA, 2010. ACM.

[31] J. Y. Tsai, P. G. Kelley, P. H. Drielsma, L. F. Cranor, J. I. Hong, and N. M. Sadeh. Who's viewed you?: the impact of feedback in a mobile location-sharing application. In D. R. O. Jr., R. B. Arthur, K. Hinckley, M. R. Morris, S. E. Hudson, and S. Greenberg, editors, *CHI*, pages 2003–2012. ACM, 2009.

[32] L. Wang, D. Wijesekera, and S. Jajodia. A logic-based framework for attribute based access control. In *Proceedings of the Second ACM Workshop on Formal Methods in Security Engineering (FMSE 2004)*, pages 45–55, Oct. 2004.

[33] J. Wickramasuriya, M. Datt, S. Mehrotra, and N. Venkatasubramanian. Privacy protecting data collection in media spaces. In *Proceedings of the 12th Annual ACM International Conference on Multimedia*, pages 48–55, New York, NY, USA, Oct. 2004.

[34] W. H. Winsborough, K. E. Seamons, and V. E. Jones. Automated trust negotiation. In *DARPA Information Survivability Conference and Exposition*, Jan. 2000.

[35] B. Wongchaowart and A. J. Lee. Oblivious enforcement of hidden information release policies. In *Proceedings of the Fifth ACM Symposium on Information, Computer, and Communications Security (ASIACCS)*, Apr. 2010.

APPENDIX

A. POINT-LOSS FUNCTIONS

This section contains the curves used for determining the point-loss functions. Figure 9(a) shows the curves we fitted to the results of the first study for individual queriers. Figure 10(a) shows the fitted curve for the case of aggregate number of accesses. Both curves use an equation as indicated in Equation 2.

From these figures we then extracted the probability functions, as shown in Figure 9(b) for the normal case and Figure 10(b) for the case of aggregate accesses.

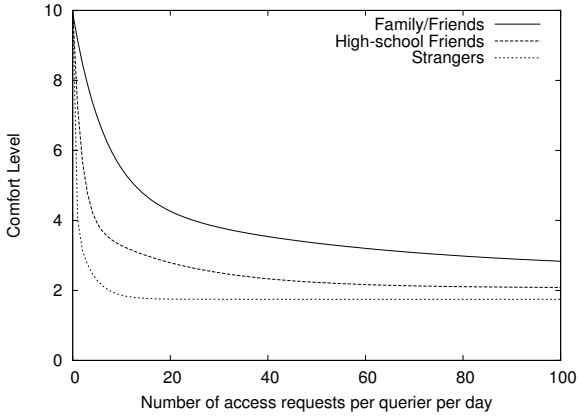

(a) This figure shows the curves fitted for the results from study one for different categories of queriers. The comfort level versus the number of access requests per querier per day is distinctly different for the three categories.

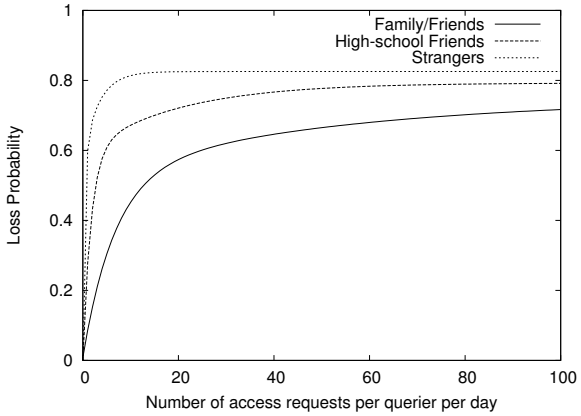

(b) This figure shows the point loss probabilities versus number of access requests per day from different categories of queriers. The loss probability increases much more sharply for strangers and high-school friends than for family/friends.

Figure 9: This figure shows the fitted curves and point loss probability functions for different categories of people.

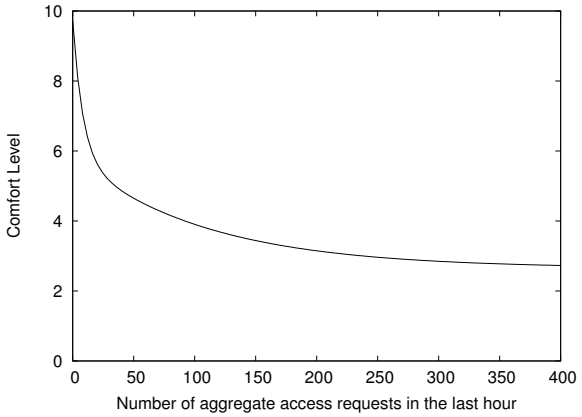

(a) This figure shows the curve fitted through the results of comfort level versus aggregate accesses in the last hour.

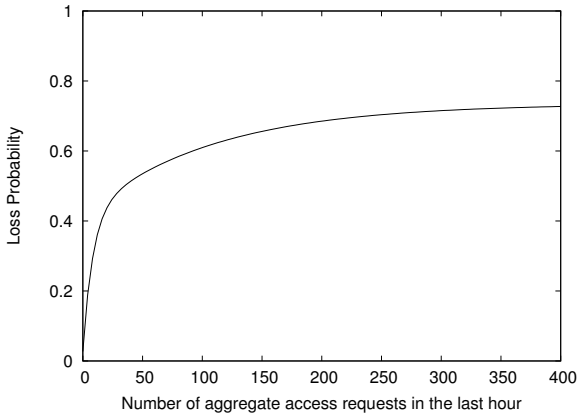

(b) This figure shows the point loss probability versus number of aggregate access requests of all queriers in the last hour. The probability of losing points increases sharply for the first 50 or so accesses and then flattens off a bit.

Figure 10: This figure shows the fitted curve and the point loss probability function for the case of aggregated accesses.